



## UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
[www.uspto.gov](http://www.uspto.gov)

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/824,068	04/13/2004	Luis Felipe Cabrera	MS128810.03/40062.139USC2	7004
23552	7590	06/25/2008	EXAMINER	
MERCHANT & GOULD PC			KERZINGER, ALEKSANDR	
P.O. BOX 2903			ART UNIT	PAPER NUMBER
MINNEAPOLIS, MN 55402-0903			2162	
MAIL DATE		DELIVERY MODE		
06/25/2008		PAPER		

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

<b>Office Action Summary</b>	<b>Application No.</b> 10/824,068	<b>Applicant(s)</b> CABRERA ET AL.
	<b>Examiner</b> ALEKSANDR KERZHNER	<b>Art Unit</b> 2162

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --  
**Period for Reply**

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
  - If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
  - Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED. (35 U.S.C. § 133).
- Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

#### Status

- 1) Responsive to communication(s) filed on 26 February 2008.  
 2a) This action is FINAL.      2b) This action is non-final.  
 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

#### Disposition of Claims

- 4) Claim(s) 21-53 is/are pending in the application.  
 4a) Of the above claim(s) \_\_\_\_\_ is/are withdrawn from consideration.  
 5) Claim(s) \_\_\_\_\_ is/are allowed.  
 6) Claim(s) 21-53 is/are rejected.  
 7) Claim(s) \_\_\_\_\_ is/are objected to.  
 8) Claim(s) \_\_\_\_\_ are subject to restriction and/or election requirement.

#### Application Papers

- 9) The specification is objected to by the Examiner.  
 10) The drawing(s) filed on 13 April 2004 is/are: a) accepted or b) objected to by the Examiner.  
     Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).  
     Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).  
 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

#### Priority under 35 U.S.C. § 119

- 12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).  
 a) All    b) Some \* c) None of:  
 1. Certified copies of the priority documents have been received.  
 2. Certified copies of the priority documents have been received in Application No. \_\_\_\_\_.  
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

\* See the attached detailed Office action for a list of the certified copies not received.

#### Attachment(s)

- 1) Notice of References Cited (PTO-892)  
 2) Notice of Draftsperson's Patent Drawing Review (PTO-948)  
 3) Information Disclosure Statement(s) (PTO-166/08)  
     Paper No(s)/Mail Date \_\_\_\_\_
- 4) Interview Summary (PTO-413)  
     Paper No(s)/Mail Date \_\_\_\_\_
- 5) Notice of Informal Patent Application  
 6) Other: \_\_\_\_\_

#### **DETAILED ACTION**

1. This office action has been issued in response to amendment filed 02/26/2008. Claims 21-53 are pending. Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action.

Accordingly, this action has been made **FINAL**, as necessitated by amendment.

#### ***Terminal Disclaimer***

2. The terminal disclaimer filed on 07/21/2006 disclaiming the terminal portion of any patent granted on this application which would extend beyond the expiration date of the full statutory term of U.S. Patent Nos. 6,553,387 and 6,735,603 has been reviewed and accepted on 08/04/2006.

#### ***Claim Rejections - 35 USC § 103***

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

3. **Claims 21-53** are rejected under 35 U.S.C. 103(a) as being unpatentable over Mann et al., US Patent No. 5,996,089 (Hereinafter "Mann et al.") in view of Shinkai et al., "HAMFS File System" (Hereinafter "Shinkai et al.").

Regarding **claim 21**, Mann et al. shows:

A computer-readable medium storing at least one extent of a logical volume having a plurality of extents, the computer-readable medium comprising (col. 7, *lines 15-29 wherein fragment files on logical volumes are discussed*):

The epoch identifier for determining a configuration status based on a comparison of the epoch identifier (Col. 9, *lines 46-56 wherein "HomeBlocks" are compared to verify integrity of the volume, "epoch identifier" is read on any identifier as "epoch" is read as a label for the identifier not imparting any patentable weight*), wherein the configuration status indicates whether the logical volume can be exposed as on line (Col. 9, *lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed*).

Mann et al. does not expressly disclose that an epoch identifier is associated with each at least one extent and that the comparison of the epoch identifier is done from each extent of the logical volume.

However, Shinkai et al. teaches:

An epoch identifier associated with each at least one extent (*page 194-195, section "3.1 Tokens" wherein each token (reading on identifier, "epoch identifier" is read on any identifier as "epoch" is read as a label for the identifier not imparting any patentable weight) is associated with at least one extent*); and

The comparison of the epoch identifier is done from each extent of the logical volume. (*page 194-195, section "3.1 Tokens" wherein each token (reading on identifier, "epoch identifier" is read on any identifier as "epoch" is read as a label for the identifier*

*(not imparting any patentable weight) is associated with at least one extent, when tokens are compared they are compared for each extent involved).*

Mann et al. teaches RAID storage system that allocated fragmented files in distributed computer system. Mann et al. performs data integrity checks on the level of a volume.

Shinkai et al. teaches the use of tokens on the level of extents in order to provide for data integrity.

It would have been obvious to one of ordinary skill in the art at the time the invention was made to provide integrity checks on the level of a file extent as taught by Shinkai et al. to supplement or replace the integrity check of data volume of Mann et al. for the predictable result of providing finer scale integrity checks.

Regarding **claim 22**, Mann et al. in view of Shinkai et al. shows:

The epoch identifier further comprises: a cluster system identifier (*Mann et al.: Fig. 6#4, Col. 9, line 67, Col. 10, lines 1-3; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 23**, Mann et al. in view of Shinkai et al. shows:

The epoch identifier further comprises: a logical volume identifier (*Mann et al.: Col. 9, lines 50-53; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 24**, Mann et al. in view of Shinkai et al. shows:

The epoch identifier further comprises: a cluster system identifier and a logical volume identifier. (*Mann et al.: Col. 9, line 67 and lines 50-53, Col. 10, lines 1-3; Shinkai*

*et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens")*

Regarding **claim 25**, Mann et al. in view of Shinkai et al. shows:

The computer-readable medium has stored thereon all of the plurality of extents for the logical volume. (*Mann et al.: Col. 9, lines 57-67; Shinkai et al.: pages 193-194, Section: "3. Architecture"*)

Regarding **claim 26**, Mann et al. in view of Shinkai et al. shows:

The epoch identifier further comprises: a number that is incremented upon each change in volume configuration. (*Mann et al.: Col. 10, lines 17-20; Shinkai et al.: page 194-195, section "3.1 Tokens" as with each change time changes and as time is changed by incrementing time is read on the number that is incremented*)

Regarding **claim 27**, Mann et al. in view of Shinkai et al. shows:

The epoch identifier further comprises: a timestamp indicating a time a volume configuration was changed. (*Mann et al.: Fig. 6D; Shinkai et al.: page 194-195, section "3.1 Tokens"*)

Regarding **claim 28**, Mann et al. in view of Shinkai et al. shows:

The data structure further comprises: an extent size; an extent identifier. (*Mann et al.: Fig. 6D, wherein specifically disclosed are data object stripe size and specific data object name in the data structure; Shinkai et al.: page 194-195, section "3.1 Tokens"*)

Regarding **claim 29**, Mann et al. in view of Shinkai et al. shows:

a computer-readable medium as claimed in claim 21 (see rejection of claim 1 above); and a processor that accesses data stored in the logical volume based on the configuration status of the logical volume determined from the comparison of the epoch identifier from each extent of the logical volume. (Mann et al.: Col. 9, lines 46-56 wherein "HomeBlocks" are compared to verify integrity of the volume, if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens," page 198, chapter "4.3 Crash Recovery")

Regarding **claim 30**, Mann et al. shows:

A data storage subsystem comprising:

A first computer-readable medium storing one or more first extents associated with a first logical volume (col. 7, lines 15-29 wherein fragment files on logical volumes are discussed), the epoch identifier for determining a configuration status of the first logical volume based on a comparison of the epoch identifier (Col. 9, lines 46-56 wherein "HomeBlocks" are compared to verify integrity of the volume, "epoch identifier" is read on any identifier as "epoch" is read as a label for the identifier not imparting any patentable weight), wherein the configuration status indicates whether the first logical volume can be exposed as on line (Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed); and

A second computer-readable medium storing one or more second extents associated with a copy of the first logical volume (col. 7, lines 15-29 wherein fragment

Art Unit: 2167

*(files on logical volumes are discussed), the copy epoch identifier for determining a configuration status of the copy of the first logical volume based on a comparison of the copy epoch identifier (Col. 9, lines 46-56 wherein "HomeBlocks" are compared to verify integrity of the volume, "epoch identifier" is read on any identifier as "epoch" is read as a label for the identifier not imparting any patentable weight).*

Mann et al. does not expressly disclose that each first or second extent includes a first data structure storing an epoch identifier and that the comparison of the epoch identifier is done from each extent of the logical volume.

However, Shinkai et al. teaches:

Each first or second extent includes a first data structure storing an epoch identifier (page 194-195, section "3.1 Tokens" wherein each token (reading on identifier, "epoch identifier" is read on any identifier as "epoch" is read as a label for the identifier not imparting any patentable weight) is associated with at least one extent); and

The comparison of the epoch identifier is done from each extent of the logical volume. (page 194-195, section "3.1 Tokens" wherein each token (reading on identifier, "epoch identifier" is read on any identifier as "epoch" is read as a label for the identifier not imparting any patentable weight) is associated with at least one extent, when tokens are compared they are compared for each extent involved)

Mann et al. teaches RAID storage system that allocated fragmented files in distributed computer system. Mann et al. performs data integrity checks on the level of a volume.

Shinkai et al. teaches the use of tokens on the level of extents in order to provide for data integrity.

It would have been obvious to one of ordinary skill in the art at the time the invention was made to provide integrity checks on the level of a file extent as taught by Shinkai et al. to supplement or replace the integrity check of data volume of Mann et al. for the predictable result of providing finer scale integrity checks.

Regarding **claim 31**, Mann et al. in view of Shinkai et al. shows:

A cluster service component including a third data structure storing a volume epoch identifier. (*Mann et al.: Col. 10, lines 25-27; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 32**, Mann et al. in view of Shinkai et al. shows:

Each of the epoch identifiers and the copy epoch identifiers have the same value. (*Mann et al.: Col. 9, lines 46-56 wherein "HomeBlocks" are compared to verify integrity of the volume; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 33**, Mann et al. in view of Shinkai et al. shows:

The cluster service component may be set by a user to expose the first logical volume as on line if the epoch identifiers of each first extent of the first logical volume are the same as the volume epoch identifier. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, and Col. 10, lines 49-62; Shinkai et al.: pages 193-194,*

*Section: "3. Architecture" and page 194-195, section "3.1 Tokens," page 198, chapter "4.3 Crash Recovery")*

Regarding **claim 34**, Mann et al. in view of Shinkai et al. shows:

The cluster service component may be set by a user to expose the first logical volume as on line only if the epoch identifiers of each first extent associated with the first logical volume and the copy epoch identifiers of each second extent of the copy of the first logical volume are the same as the volume epoch identifier. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, and Col. 10, lines 49-62; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens," page 198, chapter "4.3 Crash Recovery")*

Regarding **claim 35**, Mann et al. in view of Shinkai et al. shows:

Each first data structure includes a cluster system identifier; and each second data structure includes the cluster system identifier. (*Mann et al.: Col. 9, line 67 and lines 50-53, Col. 10, lines 1-3; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens")*

Regarding **claim 36**, Mann et al. in view of Shinkai et al. shows:

Each first data structure includes a first logical volume identifier; and each second data structure includes a second logical volume identifier. (*Mann et al.: Col. 9, lines 50-53; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens")*

Regarding **claim 37**, Mann et al. in view of Shinkai et al. shows:

The first and second logical volume identifiers are the same. (*Mann et al.: Col. 9, lines 50-59*)

Regarding **claim 38**, Mann et al. shows:

A computer-readable medium having computer-executable instructions for performing a method for exposing a logical volume as on line, the method comprising:

Reading an epoch value from the logical volume; (*col. 7, lines 15-29 wherein fragment files on logical volumes are discussed; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight*),

Reading a copy epoch value from a mirrored copy of the logical volume; (*col. 7, lines 15-29 wherein fragment files on logical volumes are discussed*),

Comparing the epoch value from each logical volume and the copy epoch value from a mirrored copy of the logical volume; (*Col. 9, lines 46-56 wherein "HomeBlocks" are compared to verify integrity of the volume; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight*)

Receiving a user selection indicating a selected consistency level; (*Col. 9, lines 24-32 wherein RAID controller can be configured by a user to be RAID-0, RAID-5, RAID-6 or RAID-7*) and

Determining a configuration status based on the comparison of the epoch value from the logical volume and the copy epoch value from a mirrored copy of the logical volume, and the selected consistency level, wherein the configuration status indicates whether the first logical volume can be exposed as on line. (*Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless*

*recovery operation is performed, recovery is based on RAID level; “epoch value” is read on any value as “epoch” is read as a label for the value not imparting any patentable weight)*

Mann et al. does not expressly disclose that

Epoch values used are from each extent of the volume.

However, Shinkai et al. shows:

*Epoch values are associated with each extent. (page 194-195, section “3.1 Tokens” wherein each token (reading on identifier, “epoch value” is read on any value as “epoch” is read as a label for the value not imparting any patentable weight) is associated with extents)*

Mann et al. teaches RAID storage system that allocated fragmented files in distributed computer system. Mann et al. performs data integrity checks on the level of a volume.

Shinkai et al. teaches the use of tokens on the level of extents in order to provide for data integrity.

It would have been obvious to one of ordinary skill in the art at the time the invention was made to provide integrity checks on the level of a file extent as taught by Shinkai et al. to supplement or replace the integrity check of data volume of Mann et al. for the predictable result of providing finer scale integrity checks.

Regarding **claim 39**, Mann et al. in view of Shinkai et al. shows:

Determining a configuration status that exposes the logical volume as on line only when the epoch values and the copy epoch values are equal if the selected

consistency level is a first consistency level. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are stripped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 40**, Mann et al. in view of Shinkai et al. shows:

Determining a configuration status that exposes the logical volume as on line when the epoch value of each extent of the logical volume are equal if the selected consistency level is a second consistency level. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are stripped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 41**, Mann et al. in view of Shinkai et al. shows:

determining the configuration status that exposes the mirrored copy of the logical volume as on line when the copy epoch value of each extent of the mirrored copy of the

logical volume are equal if the selected consistency level is the second consistency level. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are stripped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 42**, Mann et al. in view of Shinkai et al. shows:

Maintaining a volume epoch value (*Mann et al.: Col. 7, lines 15-29*);

Comparing the epoch value from each extent of the logical volume, the copy epoch value from each extent of a mirrored copy of the logical volume and the volume epoch value (*Mann et al.: Col. 9, lines 46-56 wherein "HomeBlocks" are compared to verify integrity of the volume; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight*); and

Determining a configuration status based on the comparison of the epoch value from each extent of the logical volume, the copy epoch value from each extent of a mirrored copy of the logical volume and the volume epoch value, and the selected consistency level. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock"*)

*would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are striped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens")*

Regarding **claim 43**, Mann et al. in view of Shinkai et al. shows:

*Determining a configuration status that exposes the logical volume as on line only when the epoch values, the copy epoch values, and the volume epoch value are all equal if the selected consistency level is a first consistency level. (Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are striped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens")*

Regarding **claim 44**, Mann et al. in view of Shinkai et al. shows:

*Determining comprises: determining a configuration status that exposes the logical volume as on line when the epoch value of each extent of the logical volume and the volume epoch value are equal if the selected consistency level is a third consistency level. (Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is*

Art Unit: 2167

*based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are striped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens")*

Regarding **claim 45**, Mann et al. in view of Shinkai et al. shows:

*Determining the configuration status that exposes the mirrored copy or the logical volume as on line when the copy epoch value of each extent of the mirrored copy of the logical volume and the volume epoch value are equal if the selected consistency level is the third consistency level. (Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are striped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens")*

Regarding **claim 46**, Mann et al. shows:

A computer-implemented method for exposing a logical volume as on line, the method comprising:

Reading an epoch value from the logical volume (*col. 7, lines 15-29 wherein fragment files on logical volumes are discussed; “epoch value” is read on any value as “epoch” is read as a label for the value not imparting any patentable weight*);

Reading a copy epoch value from a mirrored copy of the logical volume (*col. 7, lines 15-29 wherein fragment files on logical volumes are discussed*); and

Determining a configuration status based on a comparison of the epoch value from the logical volume, the copy epoch value from a mirrored copy of the logical volume, and a selected consistency level (*Col. 9, lines 24-32 wherein RAID controller can be configured by a user to be RAID-0, RAID-5, RAID-6 or RAID-7*), wherein the configuration status indicates whether the first logical volume can be exposed as on line. (*Col. 9, lines 46-56 wherein if “HomeBlock” does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level; “epoch value” is read on any value as “epoch” is read as a label for the value not imparting any patentable weight*)

Mann et al. does not expressly disclose that

Epoch values used are from each extent of the volume.

However, Shinkai et al. shows:

Epoch values are associated with each extent. (*page 194-195, section “3.1 Tokens” wherein each token (reading on identifier, “epoch value” is read on any value as “epoch” is read as a label for the value not imparting any patentable weight) is associated with extents*)

Mann et al. teaches RAID storage system that allocated fragmented files in distributed computer system. Mann et al. performs data integrity checks on the level of a volume.

Shinkai et al. teaches the use of tokens on the level of extents in order to provide for data integrity.

It would have been obvious to one of ordinary skill in the art at the time the invention was made to provide integrity checks on the level of a file extent as taught by Shinkai et al. to supplement or replace the integrity check of data volume of Mann et al. for the predictable result of providing finer scale integrity checks.

Regarding **claim 47**, Mann et al. in view of Shinkai et al. shows:

Determining a configuration status that exposes the logical volume as on line only when the epoch values and the copy epoch values are equal if the selected consistency level is a first consistency level. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are stripped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 48**, Mann et al. in view of Shinkai et al. shows:

Determining a configuration status that exposes the logical volume as on line when the epoch value of each extent of the logical volume are equal if the selected consistency level is a second consistency level. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are stripped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 49**, Mann et al. in view of Shinkai et al. shows:

Determining the configuration status that exposes the mirrored copy of the logical volume as on line when the copy epoch value of each extent of the mirrored copy of the logical volume are equal if the selected consistency level is the second consistency level. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are stripped and comparison would be made on the volume itself, "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 50**, Mann et al. in view of Shinkai et al. shows:

Maintaining a volume epoch value (*Mann et al.: Col. 7, lines 15-29*); and

Determining a configuration status based on a comparison of the epoch value from each extent of the logical volume, the copy epoch value from each extent of a mirrored copy of the logical volume, the selected consistency level, and the volume epoch value. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are striped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens"*)

Regarding **claim 51**, Mann et al. in view of Shinkai et al. shows:

Determining a configuration status that exposes the logical volume as on line only when the epoch values, the copy epoch values, and the volume epoch value are all equal if the selected consistency level is the first consistency level.

Regarding **claim 52**, Mann et al. in view of Shinkai et al. shows:

Determining a configuration status that exposes the logical volume as on line when the epoch value of each extent of the logical volume and the volume epoch value are equal if the selected consistency level is a third consistency level. (*Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed*

*as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are stripped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens")*

Regarding **claim 53**, Mann et al. in view of Shinkai et al. shows:

*Determining the configuration status that exposes the mirrored copy of the logical volume as on line when the copy epoch value of each extent of the mirrored copy of the logical volume and the volume epoch value are equal if the selected consistency level is the third consistency level. (Mann et al.: Col. 9, lines 46-56 wherein if "HomeBlock" does not match a volume cannot fully be exposed as on line unless recovery operation is performed, recovery is based on RAID level, if for example RAID level is 5 "HomeBlock" would be compared on a logical volume as well as on the mirror, On the other hand if RAID is 0 the disks are stripped and comparison would be made on the volume itself; "epoch value" is read on any value as "epoch" is read as a label for the value not imparting any patentable weight; Shinkai et al.: pages 193-194, Section: "3. Architecture" and page 194-195, section "3.1 Tokens")*

#### ***Response to Arguments***

4. Applicant's arguments, see pages 9-11, filed 10/05/2007, with respect to claim rejections under 35 U.S.C. 101 and double patenting rejections of claims under judicially

Art Unit: 2167

created doctrine (in view of filed and approved terminal disclaimer) have been fully considered and are persuasive. The 35 U.S.C. 101 and double patenting rejections of Office action mailed on 04/21/2006 has been withdrawn.

5. Applicant's arguments with respect to claims 21-53 have been considered but are moot in view of the new ground(s) of rejection.

***Conclusion***

Applicant's amendment necessitated the new ground(s) of rejection presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**. See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to ALEKSANDR KERZHNER whose telephone number is (571)270-1760. The examiner can normally be reached on Mon-Fri 9:00-6:30.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John E. Breene can be reached on (571) 272-4107. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Aleksandr Kerzhner/  
Examiner, Art Unit 2162  
06/03/2008

/Kuen S Lu/  
Primary Examiner, Art Unit 2167